International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Series G
Supplement 74
(12/2021)

SERIES G: TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS

Network slicing in a passive optical network context

ITU-T G-series Recommendations – Supplement 74

ITU-T G-SERIES RECOMMENDATIONS

TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS

For further details, please refer to the list of ITU-T Recommendations.

Supplement 74 to ITU-T G-series Recommendations

Network slicing in a passive optical network context

Summary

Supplement 74 to ITU-T G-series Recommendations considers the possible implications of network slicing and related functionality (e.g., as defined in 3GPP and IMT-2020 as part of 5G wireless) in the setting of optical access networks, in particular passive optical network (PON) systems. Use cases are identified for slicing with a focus on PON slicing. The Supplement describes the functional expectations of a PON system to support PON slicing and considers the architecture for the control of PON slices. Also, the aim is to provide a clear demarcation of PON slicing such that it can be considered independently of slicing in any other segments of the network.

History

Edition	Recommendation	Approval	Study Group	Unique ID*
1.0	ITU-T G Suppl. 74	2021-12-17	15	11.1002/1000/14974

Keywords

PON, PON slicing, slicing.

* To access the Recommendation, type the URL http://handle.itu.int/ in the address field of your web browser, followed by the Recommendation's unique ID. For example, http://handle.itu.int/11.1002/1000/11830-en.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

This is an informative ITU-T publication. Mandatory provisions, such as those found in ITU-T Recommendations, are outside the scope of this publication. This publication should only be referenced bibliographically in ITU-T Recommendations.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this publication may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the publication development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents/software copyrights, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the appropriate ITU-T databases available via the ITU-T website at http://www.itu.int/ITU-T/ipr/.

Table of Contents

Supplement 74 to ITU-T G-series Recommendations

Network slicing in a passive optical network context

1 Scope

This Supplement considers the possible implications of network slicing and related functionality (e.g., as defined in 3GPP and IMT-2020 as part of 5G wireless) in the setting of optical access networks, in particular PON systems. Use cases are identified for slicing with focus on PON slicing. The Supplement describes the functional expectations of a PON system to support PON slicing and considers the architecture for the control of PON slices. Also, the aim is to provide a clear demarcation of PON slicing such that it can be considered independently of slicing in any other segments of the network.

In this Supplement, the terms "slice" and "slicing" refer to "PON slice" and "PON slicing" unless otherwise indicated.

2 References

[ITU-T G.7702] Recommendation ITU-T G.7702 (2018), *Architecture for SDN control of transport networks*.

[ITU-T G.9804.2] Recommendation G.9804.2 (2021), *Higher speed passive optical networks – Common transmission convergence layer specification.*

[ITU-T G.9807.1] Recommendation ITU-T G.9807.1 (2016), *10-Gigabit-capable symmetric passive optical network (XGS-PON).*

[ITU-T Y.3100] Recommendation ITU-T Y.3100 (2017), *Terms and definitions for IMT-2020 network*.

[ITU-T Y.3150] Recommendation ITU-T Y.3150 (2020), *High-level technical characteristics of network softwarization for IMT-2020*.

[ITU-T Y.3151] Recommendation ITU-T Y.3151 (2019), *High-level technical characteristics of network softwarization for IMT-2020 – Part: SDN*.

[BBF TR-370] BBF TR-370 Issue-2 (2020), *Fixed Access Network Sharing – Architecture and Nodal Requirements*.

3 Definitions

3.1 Terms defined elsewhere

This Supplement uses the following term defined elsewhere:

3.1.1 saturation of a flow (clause C.7.3.3 of [ITU-T G.9807.1]): A flow is **saturated** when it has been assigned its saturation level. A flow is **unsaturated** when it has been assigned less than its saturation level. The saturation level for upstream flows is defined in the PON TC Recommendations.

3.2 Terms defined in this Supplement

This Supplement defines the following terms:

3.2.1 OLT slice: An optical line termination (OLT) slice is one partition of the traffic management functions of the OLT intended to facilitate network slicing over a PON-based access network. An OLT slice might have its own management of traffic and other parameters and appear as an independent logical OLT.

3.2.2 ONU slice: An optical networking unit (ONU) slice is one partition of the traffic management and buffering functions of the ONU intended to facilitate network slicing over PON-based access network. An ONU slice might have its own management of traffic and other parameters.

3.2.3 PON slice: A PON slice is a group of one or more flows associated with one or more ONUs that are treated as a single entity by a hierarchical traffic scheduler.

3.2.4 PON virtualization: An overall approach of designing, implementing, deploying, managing and/or maintaining PON equipment and/or PON components by software programming.

3.2.5 saturation of a slice: This notion is based on the existing notion of saturation at flow level and is used for the description of the handling of upstream bandwidth by the dynamic host configuration protocol (DBA).

A slice is saturated when it either has been allocated its max configured bandwidth **or** when all of its flows are saturated. Note that in the former case one or several flows of the slice may remain unsaturated.

A slice is unsaturated when it has been allocated less than its max configured bandwidth **and** has at least one unsaturated flow.

Example cases:

Assuming a slice with max upstream bandwidth (BW) = 5 Gbit/s, that contains a collection of flows of equal priority, and each has an equal max BW = 1 Gbit/s.

Case 1) DBA allocates a slice BW equal to the max BW = 5 Gbit/s. The slice is **saturated**. This is independent of the saturation status of its constituent flows:

- e.g., if the slice contains 10 flows of equal demand for BW = 0.5 Gbit/s, then each flow can be allocated its demand by DBA and is saturated.
- e.g., if the slice contains 10 flows of equal demand for BW > 0.5 Gbit/s, then each flow can only be granted BW = 0.5 Gbit/s and is unsaturated.
- e.g., the slice contains 5 flows of equal demand for BW = 0.4 Gbit/s and 5 flows of equal demand for BW = 1 Gbit/s. The first five flows are allocated BW = 0.4 Gbit/s and are saturated, the last five flows are allocated BW = 0.6 Gbit/s and are unsaturated.

Case 2) DBA allocates a slice with BW = 2 Gbit/s (below its max BW), and each flow is allocated as much as its demand and is hence saturated (e.g., 10 flows with demand for BW = 0.2 Gbit/s and allocation of BW = 0.2 Gbit/s).

The slice is **saturated**.

Case 3) DBA allocates a slice BW of 4 Gbit/s (below its max) and

- Flows 1 to 9 each have a demand of 0.4 Gbit/s and are allocated BW = 0.4 Gbit/s by DBA and are saturated, and
- Flow 10 has a demand of 0.8 Gbit/s and can only be allocated BW = 0.4 Gbit/s by DBA and is unsaturated.

The slice is **unsaturated**.

3.2.6 slice isolation: On a PON, the ability to define and guarantee some per-slice and per-flow parameters for a given slice, regardless of the behaviour in the other slices (configuration and load), and within the limits of the shared resource.

3.2.7 surplus bandwidth for a slice: The portion of the capacity that remains available after the guaranteed bandwidth components have been dynamically assigned to the considered flows.

When there is no slice awareness, upstream surplus bandwidth is used at the channel termination (CT) level (e.g., as described in clause C.7.3.1 of [ITU-T G.9807.1]) to share non-guaranteed bandwidth between flows of the whole PON (the upstream surplus bandwidth of the channel termination is available to the flows on the channel termination).

The same concept can also be applied at slice level to share non-guaranteed bandwidth between slices, and to share per-slice non-guaranteed bandwidth between flows inside the slice (the upstream surplus bandwidth of the channel termination is available to the slices on the channel termination, the upstream surplus bandwidth of the slice is available to the flows inside the slice).

4 Abbreviations and acronyms

This Supplement uses the following abbreviations and acronyms:

BW Bandwidth

CLI Command Line Interface

CMUD Create, Modify, Update and Delete

CT Channel Termination

DBA Dynamic Bandwidth Assignment

DHCP Dynamic Host Configuration Protocol

EMS Element Management System

FA Factory Automation

FTP File Transfer Protocol

HSP High Speed PON

IGMP Internet Group Management Protocol

IP Internet Protocol

ITS Intelligent Transport System

MAC Media Access Control

MEC Mobile Edge Computing

NTP Network Time Protocol

OLT Optical Line Termination

ONU Optical Networking Unit

OSS Operation Support System

OTN Optical Transport Networking

PA Process Automation

PDH Plesiochronous Digital Hierarchy

PON Passive Optical Network

SDH Synchronous Digital Hierarchy

SDN Software Defined Network

SLA Service Level Agreement

SNI System Network Interface

SNMP Simple Network Management Protocol

SONET Synchronous Optical Networking

TBDT	Bandwidth Assignment Delay Tolerance
TC	Transmission Convergence
T-CONT	Transmission Container
TJT	Time Jitter Tolerance
TPST	Protection Switching Delay Tolerance
UE	User Equipment
UNI	User Network Interface
URLLC	Ultra-Reliable Low Latency Communication
VLAN	Virtual LAN
VN	Virtual Network
VNO	Virtual Network Operator
VPN	Virtual Private Network
VR	Virtual Reality

5 Conventions

None.

6 Slicing in optical access

6.1 Introduction

A conventional PON system is capable of supporting multiple services with associated service segregation, QoS control, and connection configuration. The existing PON Recommendations defining the transmission convergence (TC) layer (e.g., [ITU-T G.9804.2]) support the management of traffic on an OLT CT at flow level via a single control function. Slicing on OLT CTs (PON slicing) collect individual flows into groups called slices and adds a control function for sharing common resources at the slice level. This control function enables a degree of isolation of services or virtual network operators (VNO) associated with different slices when the PON resources are shared between the slices.

A PON slice groups at least several flows on one or multiple ONUs of the OLT CT. Each slice must be handled accordingly by the DBA and the hierarchical scheduler in the OLT. A PON slice is characterized by a set of parameters, i.e., a slice profile, which are for further study.

Slicing also implies coordination between management instances for the OLT, and for the ONUs in cases where ONU features would be used to handle flows at both flow and slice level. These aspects can be seen from some use cases, but these details are outside the scope of the current version of this Supplement.

6.2 Network slicing for services

Figure 6-1 illustrates a network slicing configuration for a diverse set of services requiring different quality-of-service (QoS) and thus grouped into four slice types, such as enhanced fixed broadband (eFBB), guaranteed reliable experience (GRE), Internet of things (IoT), and vertical/industrial applications. Multiple dedicated networks (D-Net's) offer services to specific market segments like residential, wholesale, smart city, industrial, hospital or any other enterprise-oriented segments. Each D-Net has a set of slice types to support. For example, the residential D-Net supports eFBB, GRE and IoT slice types to provide traditional broadband, cloud virtual reality (VR), cloud gaming, and smart home services. By sorting all the network slice instances according to their slice types, all the services can be effectively supported to meet their corresponding QoS requirements.

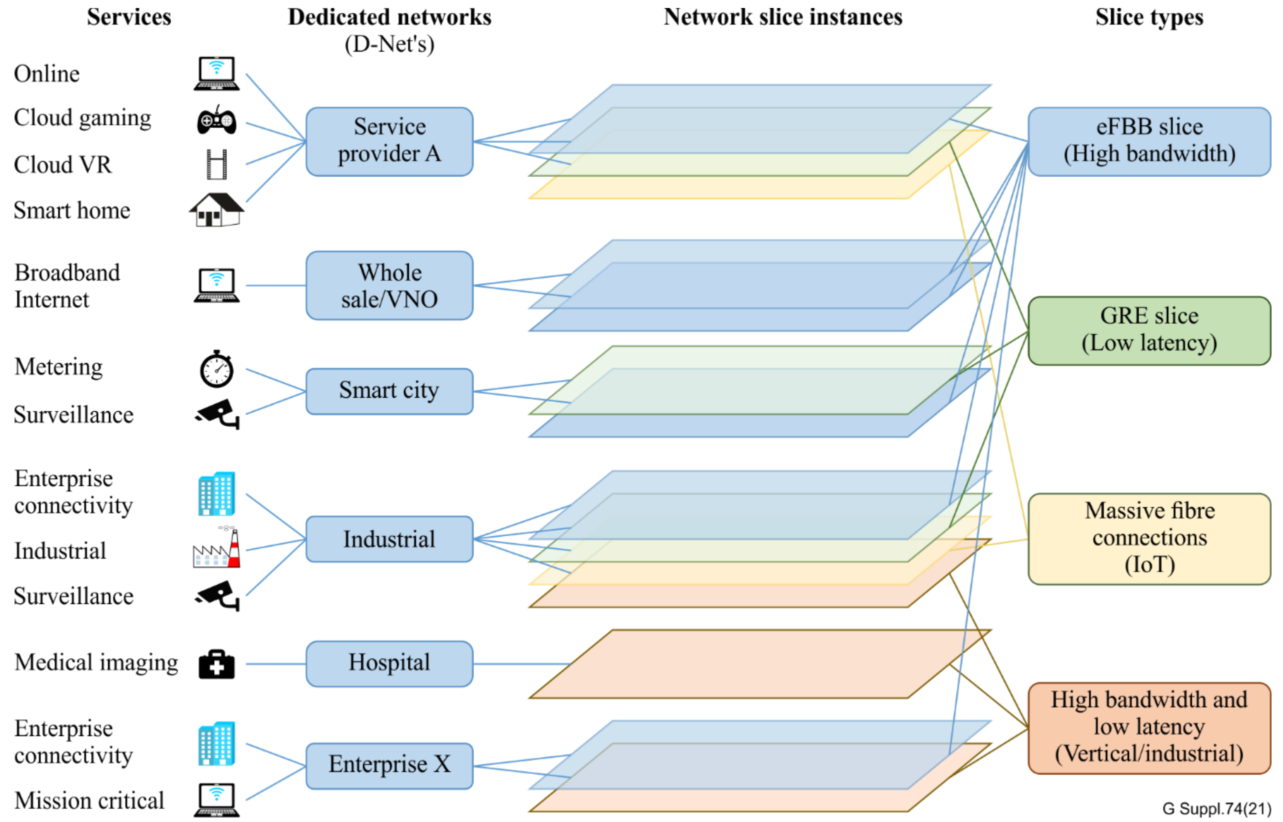

Figure 6-1 – Network slicing for services requiring different slice types (VNO: virtual network operator)

A network operator can source one or more services grouped into one or more network slices of the same or different types.

6.3 Slicing in PON systems in the context of end-to-end network slicing

In the broadest sense, the network provides end-to-end connectivity for applications running between user devices and/or application servers. The end-to-end connectivity is composed of several network segments, as illustrated in Figure 6-2, listed from right to left:

1) local network at user side (e.g., a CPN)
2) access network
3) interconnection networks (such as transport, aggregation, core) and possibly other access networks
4) networks containing application servers (e.g., data centres) or other local networks at the end-user side.

In Figure 6-2, the access network segment (2) is represented by an optical access network (OAN) based on a PON system.

NOTE – Figure 6-2 is generic and does not make a statement on where the end-points of each network segment are located and what technology is used to implement network segments other than the access network.

An end-to-end network slice spans through all the network segments between the application server and the served user devices. Each segment is characterized by a given technology and has an appropriate slicing level of its own (called "virtual networks") with the aim to maintain the integrity of the end-to-end network slicing by supporting the required network slicing parameters, specifically

QoS. A given network segment, with the aim to support the end-to-end requirements of network slicing must integrate with the slicing in its adjacent segments (i.e., provide mapping between virtual networks of adjacent segments in terms of connectivity and QoS, etc…).

Access network slicing, accordingly, integrates with virtual networks of the interconnection network at system network interfaces (SNI) and with virtual networks of the local network at user network interfaces (UNI). The following general concept of slicing is further used in PON-based access networks: in addition to the current capabilities of sharing of the PON resources via different functions, such scheduling functions are extended with an additional level in their hierarchy, to support per-slice instances of resource control (see clause 8 for more details).

The application of this slicing concept in an optical access network segment is further described in three major areas that can be considered either individually or in a combined way. As illustrated in Figure 6-2, the PON OLT equipment supports multiple OLT CTs, each OLT CT supports an ODN with multiple subtending ONUs. The areas are respectively:

(2A) OLT slicing represents slicing in the OLT part that is common to the multiple OLT CTs;

(2B) PON slicing represents an OLT CT carrying multiple PON slices to/from the PMD+TC functions of the associated ONUs;

(2C) ONU slicing represents slicing in the part of the ONU behind its PMD+TC function.

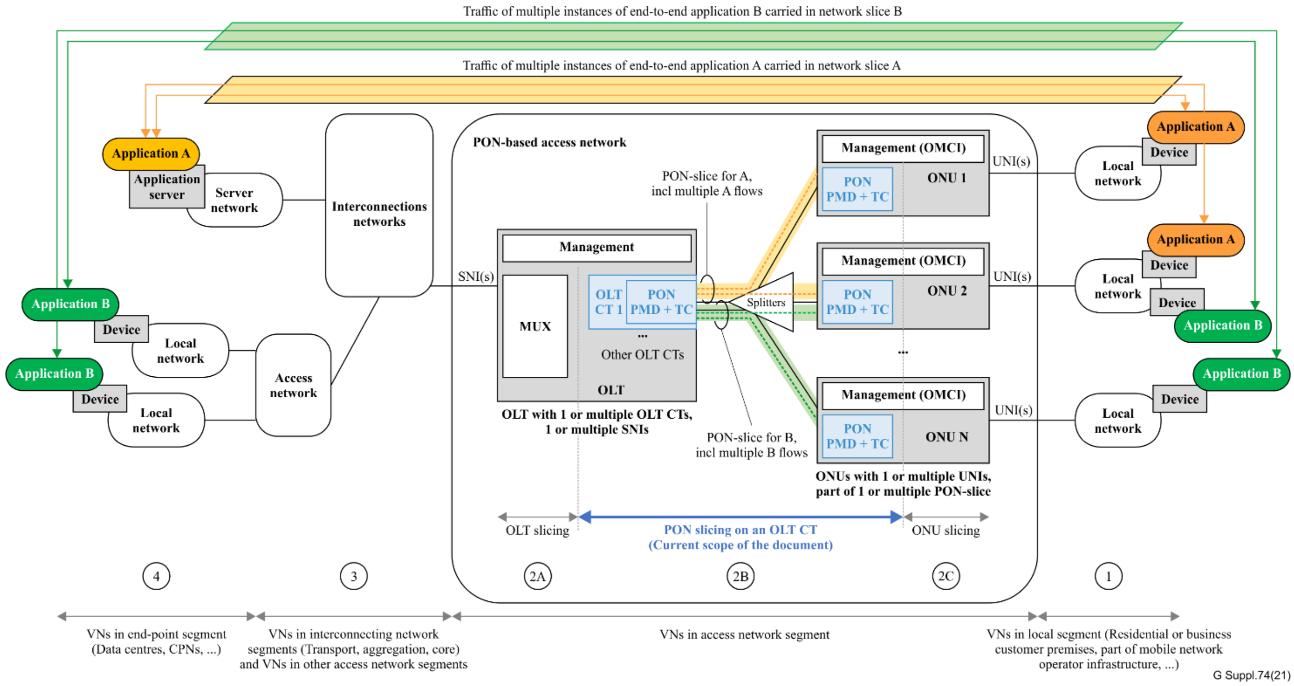

Figure 6-2 – Illustration of slicing in OAN (in particular PON slicing) in the context of network slices

It should be noted that a large number of network slices of the same and different types can be sourced by multiple service providers and VNOs. Interconnecting network segments, including optical access networks, may establish multiple slices to accommodate the appropriate transfer of each network slice. There is the possibility to map multiple virtual networks from one segment into fewer virtual networks in an adjacent segment based on the similarity of QoS requirements. In other words, a given PON slice may support multiple network slices of the same type (this is not indicated in Figure 6-2).

In a slice-aware optical access network segment, the optical line terminal (OLT) is slice-aware, while the optical network unit (ONU) can be either slice-unaware or slice-aware:

- In case all services associated with the ONU belong to one slice, the ONU can be slice-unaware.
- In case, services associated with the ONU belong to different slices, depending on QoS requirements per slice, the ONU may or may not need slice-awareness in order to support the QoS requirement for the user data associated with each slice.

In the case where the ONU is slice-unaware, the mapping/de-mapping between PON slices and the service content of network slices happens at the OLT.

In the case where the ONU is slice-aware, the mapping/de-mapping between PON slices and service content of network slices happens at both the OLT and at the ONU.

6.4 The scope of slice awareness in PON systems

The PON-based access network segment (PON system composed of OLT and ONUs) must provide seamless transport of network slices between the OLT SNI(s) and the ONU UNI(s), taking into account appropriate mappings for QoS and connectivity. This can include the following levels of slice-awareness:

PON slicing involves the partition of the shared resources and functions of the OLT CT, namely the upstream and downstream PON capacity and the control of their sharing by utilising dynamic bandwidth allocation (DBA) in upstream and hierarchical scheduling downstream. A conventional PON has only a single DBA level that implements the bandwidth fairness algorithm at the flow level as described in the ITU PON TC Recommendations (e.g., [ITU-T G.9804.2]). In a sliced PON, the DBA function of the OLT CT additionally has to share its resources at the slice level. The internal architecture or modularity to manage multiple levels and their instances is out of the scope of this Supplement. This Supplement also describes the support of PON slicing, namely the sharing of the common capacity of the PON CT by multiple PON slices. In the downstream direction, a conventional PON hierarchy in scheduling is extended with a level for slicing. The internal architecture of such a hierarchy is out of the scope of this Supplement. Clause 8 describes the mechanisms for PON slicing at a functional level.

ONU slicing involves the partition of the traffic management and buffering functions of the ONU. A conventional ONU has corresponding capabilities (e.g., traffic management) that are controlled via OMCI. A sliced ONU could have multiple similar functional instances, each controlled fairly independently. Clearly, an indeterminate coordination or unification of all these different functional instances is assumed. Clause 6.5 describes the two possible scenarios from an ONU perspective, namely slice-aware ONU, and slice-unaware ONU. Further considerations for ONU slicing are for future study.

OLT slicing partitions the relevant traffic management functions of the OLT. Similar to the ONU, the only resources that are shared in a conventional OLT (which does not support slicing), are the management and control system. Thus, a sliced OLT would appear to north-bound network management systems as several logically independent OLTs. There are many ways that the resources of the OLT associated with a particular OLT CT could be divided. OLT slicing is for future study.

The other segments (the aggregation and transport networks feeding the OLT and the local (customer's) network attached to the ONU) are out of the scope of this Supplement.

6.5 ONU scenarios in network slicing

The position of an ONU in network slicing can be classified into two general types:

1) Multiple slices per ONU: A single ONU carries multiple network slices simultaneously. Multiple ONUs may be connected to each network slice.

2) Single slice per ONU: An ONU carries only a single network slice. Multiple ONUs may be connected to that single network slice.

An OLT can connect a mixture of ONUs carrying a single slice and ONUs carrying multiple slices. In some cases, an ONU will not be dedicated or optimized to a particular slice type and can carry single or multiple slices without ONU slicing. In other cases, an ONU will be slice aware and ONU slices are optimized to one or more network slice types (e.g., for performance reasons).

Slicing scenarios include an ONU carrying multiple slices, an ONU carrying a single slice, and a mixture as shown in Figure 6-3, Figure 6-4, and Figure 6-5 below.

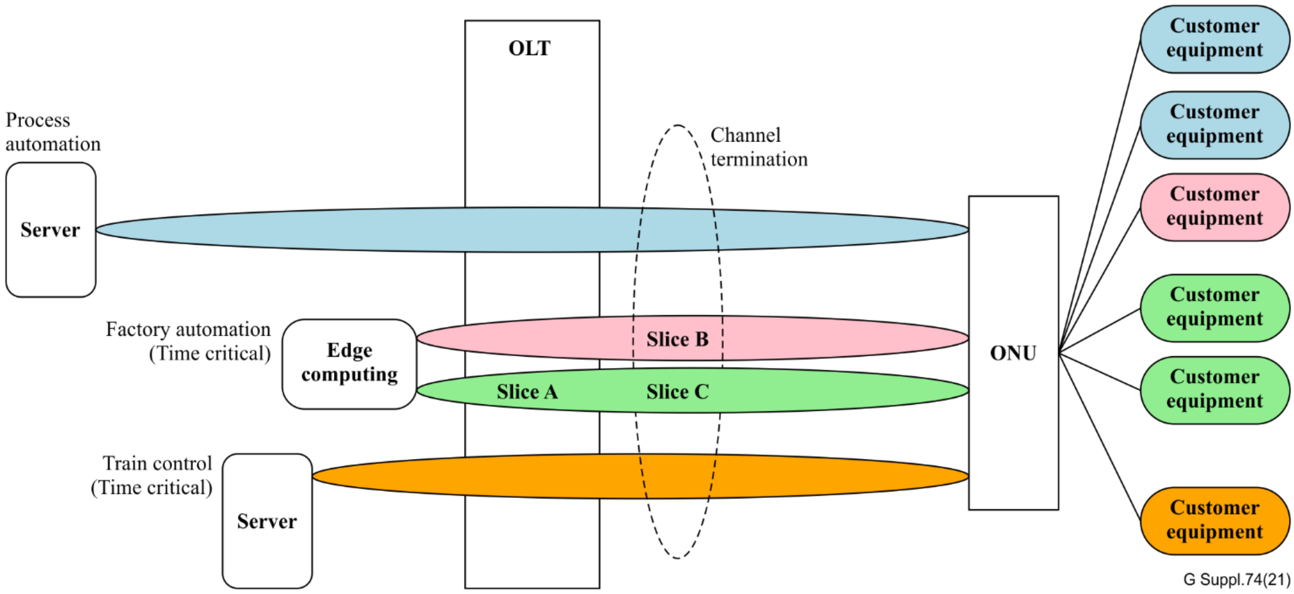

Figure 6-3 – One ONU carrying multiple slices, each slice with different requirements for rate, delay, isolation, and routing

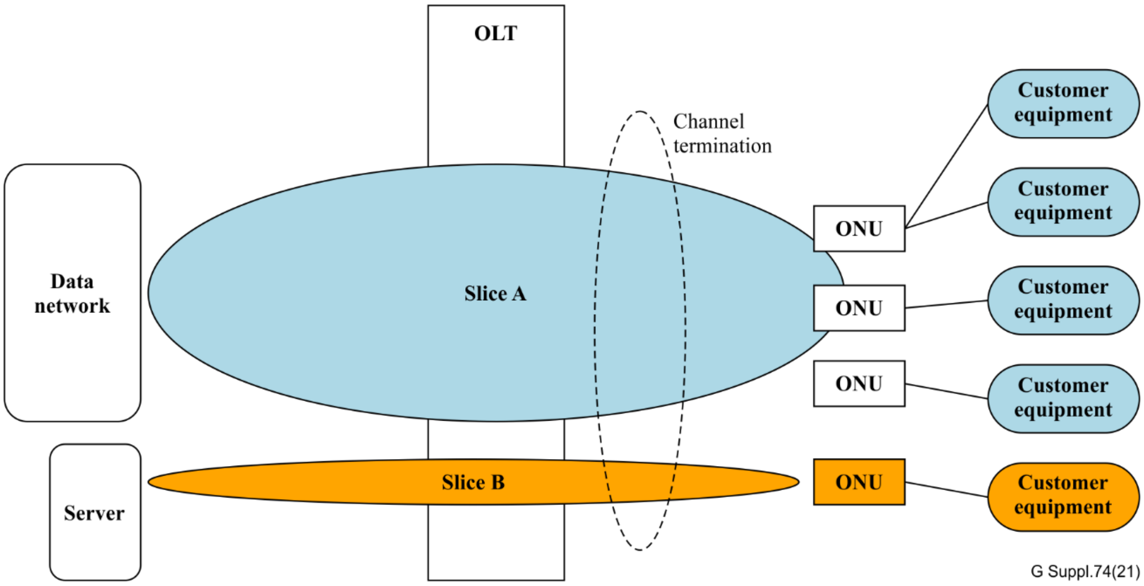

Figure 6-4 – ONU dedicated to a single slice. The ONU may be optimized to the properties of slice B

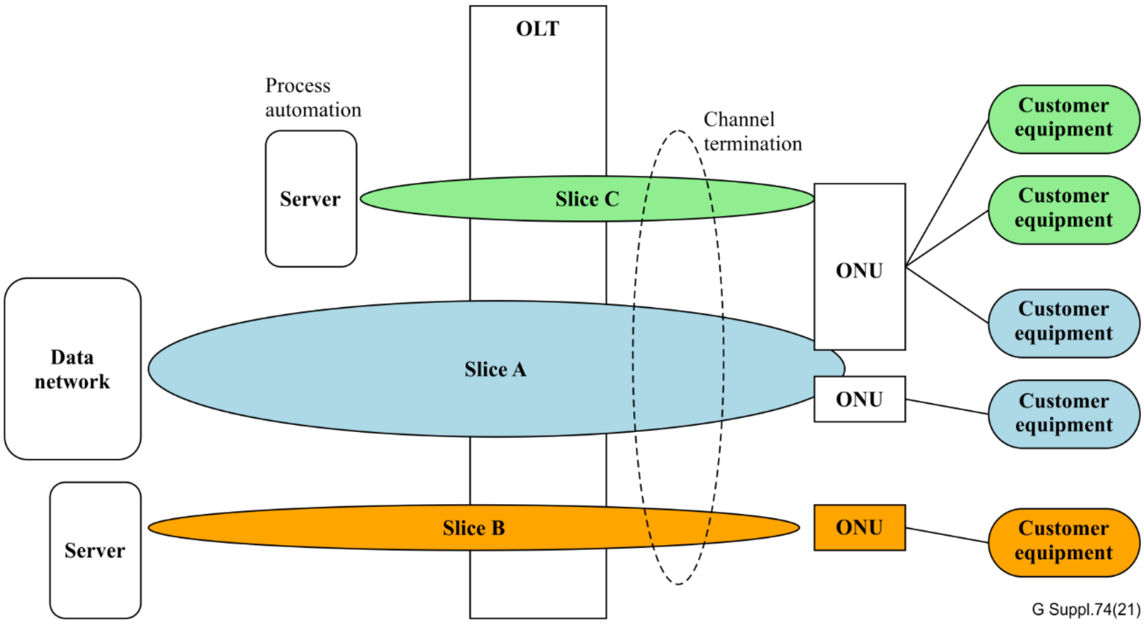

Figure 6-5 – Mix of ONUs carrying single or multiple slices

7 Use cases for slicing in PON-based optical access

7.1 Use case 1: Network slicing for PON supporting several applications of different latency constraints

When providing an application service that has a specific requirement regarding latency of data communication, a slice can set the requirement for each application and then the expected latency can be effectively provided.

Figure 7-1 is reproduced from Figure I.1 of [ITU-T Y.3151] and shows an example of applications with different latency requirements that run on the same optical access infrastructure: control of devices/systems such as process automation (PA)/factory automation (FA)/intelligent transport system (ITS) / train control (TC), mobile fronthaul or backhaul and server applications.

Server applications mainly deal with transactions for non-URLLC (non-ultra-reliable low latency communication) applications and relevant information delivery/aggregation. Non-URLLC application is an application category in which ultra-low latency is not requested. However, a certain level of latency and processing time is often required. For example, "xx ms" on straight arrows in Figure 7-1 shows transfer latency time, "xx ms" on curved arrows are the requested processing periods.

The services for FAs and ITS have the most stringent delay requirements e.g., a maximum return time trip of 40 ms, so mobile edge computing (MEC) is used in the access node to achieve the minimum delay time. Services for PAs and TCs do not use MECs but still require end-to-end delay control. Also, mobile services can carry URLLC and non-URLLC applications with an end-to-end latency of a few ms in the case of URLLC.

In Figure 7-1, slices-A/B/C/D include user equipment (UE), access nodes, servers or mobile edge computing (MEC) to serve a different kind of control of devices/systems applications. Access nodes conduct real-time resource allocation by grasping both the status and the quantity of resources for the applications to meet the requirements of transfer latency time in each network slice. The access nodes also decide which route of data flow is appropriate (i.e., a data flow to servers or the MEC).

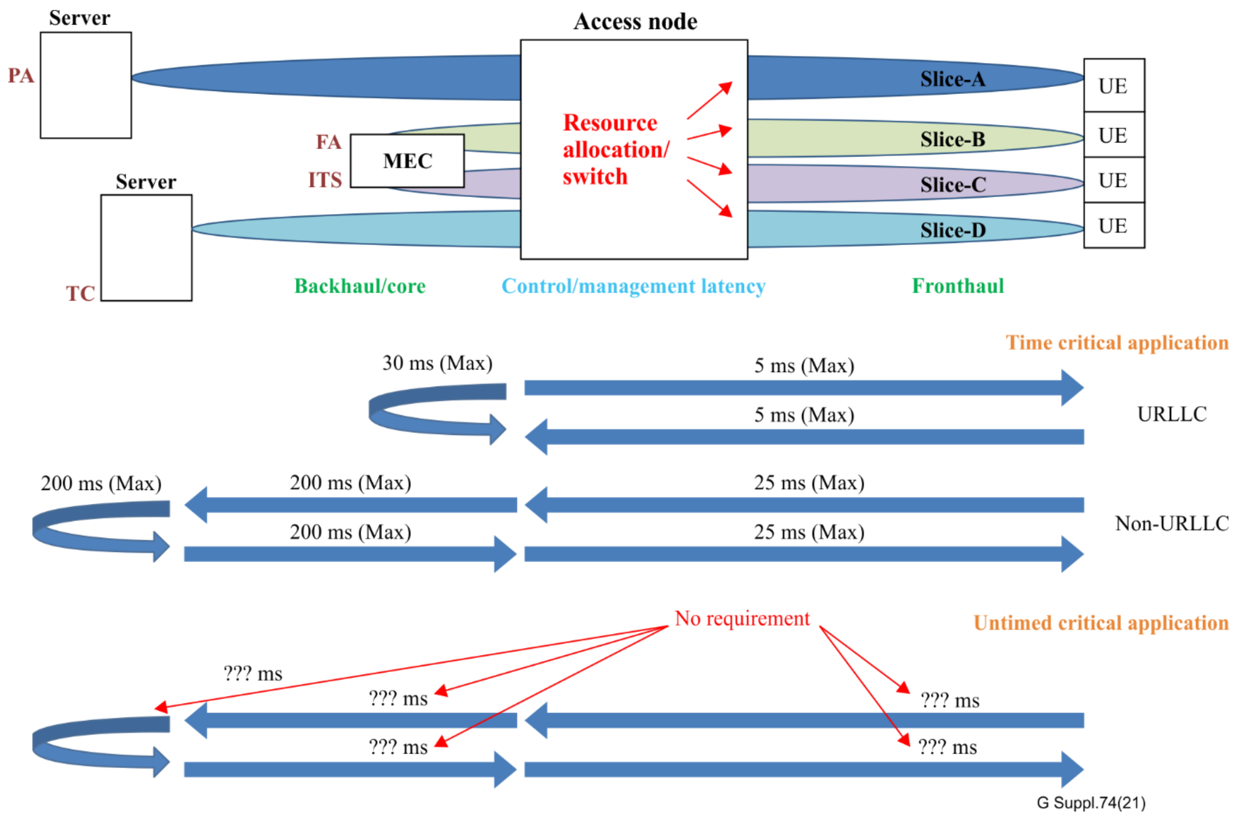

Figure 7-1 – Example of time-critical applications. Potentially, 1 ONU related to 1 user equipment (UE) can support several applications of different latency constraint

Impacts on PON slices

To achieve proper control of the transmission delay with regard to the different services, different PON slices are used. For instance, a slice serving multiple UEs, flows requesting the same delay time are grouped together, and then their delay time is controlled. Thus, various services with different delay requirements can be accommodated simultaneously on the same PON system by building PON slices.

In case several applications with different latency constraints need to be supported by a given ONU, several PON slices, with different latency requirements must be handled for this particular ONU. This has an impact on the OLT resource allocation scheme.

7.2 Use case 2: Dynamic optimization of resource allocation in a PON supporting mobile applications

This use case considers a mobile communication network that aims at optimizing resource allocation based on the amount of users / traffic per slice.

Figure 7-2 is reproduced from [ITU-T Y.3151] Figure I-2 and shows a network that consists of user equipment (UE), mobile network PON-based access network, backhaul/core network and a server. Server conducts data delivery to UEs or aggregates data from UEs and treats the relevant information processing. The network slice encompasses the PON-based segment of the network that is also sliced, i.e., consists of a group of flows in which some or all can belong to the same logical network.

When users, cars, etc. move, it is not necessary to activate all the mobile cells, and associated slices can provide the minimum resources. Within a given slice, as the user traffic is moving from cell to cell, part of the resources associated with this slice in the corresponding cell dynamically changes, and accordingly, changes the used bandwidth of the slice in different ONUs and in the PON CT.

Figure 7-2 shows the resource allocation and the possible ways for its optimization in the case of mobile applications. When there are many users in cell A, most of the resources (e.g., bandwidth) of slice-X are allocated to cell A, and less to cell B and cell C. Then, as the users move to cell B and then C, more resources are allocated to cell B and cell C but less to cell A and so on. Thus, versus time, mobile cells can be operated with the least amount of resources required. If other slices are formed in the same cell in addition to slice-X, they do not affect each other because inter-slice isolation can be achieved in the PON channel termination. This change of used traffic in mobile cells results in corresponding changes of the used bandwidth associated with slice-X of the sourcing ONUs and corresponding traffic scheduling patterns in the OLT.

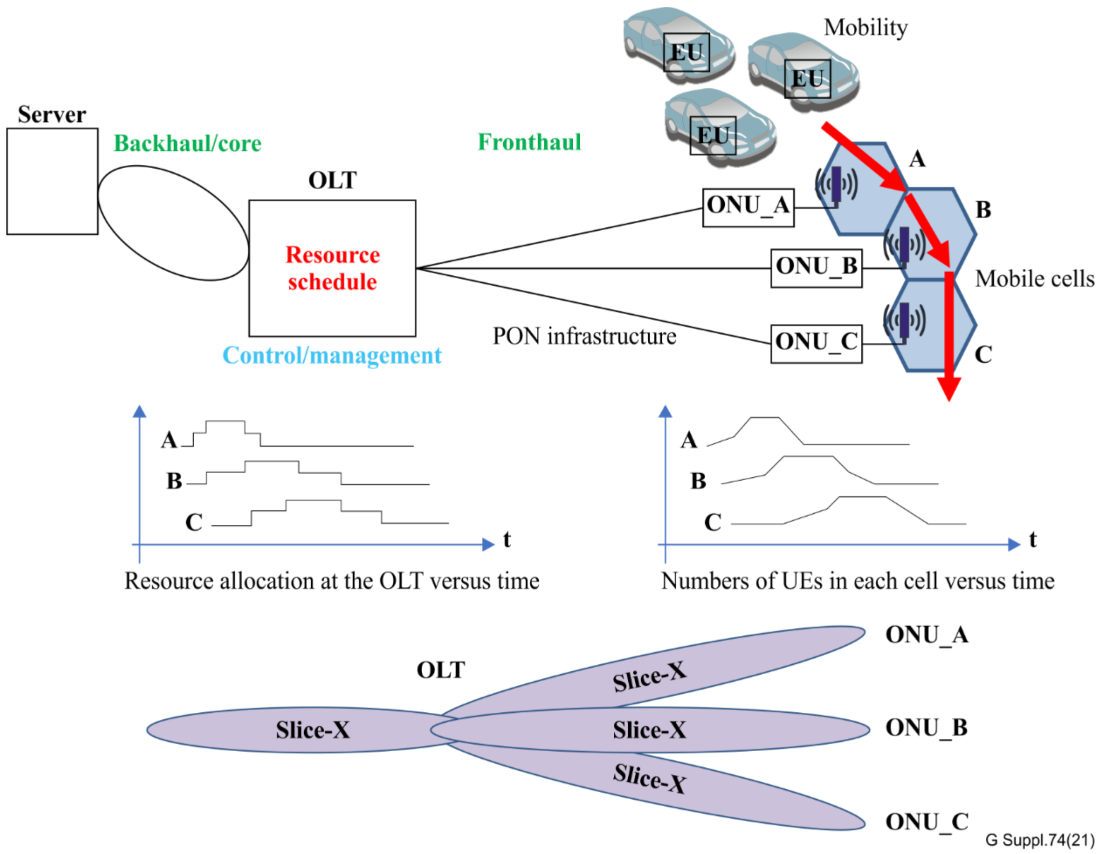

Figure 7-2 – Example of resource allocation in case of mobile application illustrating the dynamic behaviour of PON-based access supporting mobile applications

Impacts on PON slicing

The dynamic optimization of the bandwidth allocation can require PON slicing which has to be managed accordingly. Concerning slice management, two processes and two time constants can be identified.

In a fast (milliseconds) time frame, it can be required to optimize the share of the bandwidth between ONUs within a particular slice, for instance by increasing the bandwidth allocated to one ONU and reducing it to another one, and so on. This dynamic behaviour impacts the way the OLT traffic schedulers distribute the resources related to a given slice.

In a slow (seconds) time frame, a reconfiguration of slicing and slice parameters (e.g., ONUs and flows are part of the slice, associated bandwidth on the OLT CT) can be arranged. The DBA needs to be able to adapt to such reconfigurations of the PON slice.

7.3 Use case 3: Slicing in PON systems for multi-operator scenarios

In an optical access network, by logically dividing a physical OLT into multiple slicing OLTs, the key resources such as access equipment rooms, OLTs and ODNs are shared to maximize the investment value.

The main application scenarios of OLT slicing are multi-service slicing and multi-operator slicing. Multiple operator slicing is a physical OLT device which is used to provide access services for multiple operators. In [BBF TR-370] for fixed access network sharing (FANS), an operator who owns the OLT equipment is called an infrastructure operator, while an operator who shares the OLT equipment to provide service to end-users is called a virtual network operator. To maximize the sharing of infrastructure network resources, PON-level slicing is recommended as shown in Figure 7-3.

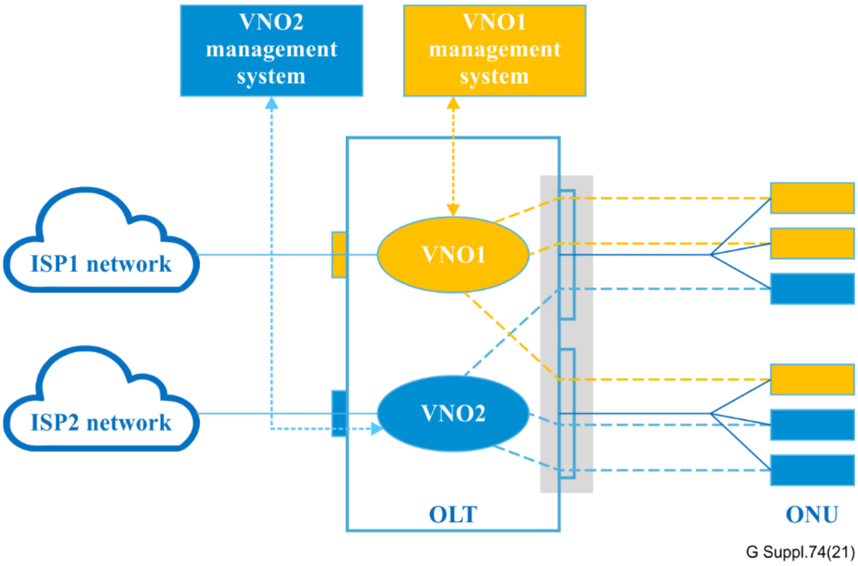

Figure 7-3 – Architecture of OLT slicing for a multi-operator

To completely isolate services among slices, the management plane, control plane, and forwarding plane of each slice need to be independent of each other.

– Management plane: Each slice has its own in-band management channel, management internet protocol (IP), system media access control (MAC), simple network management protocol (SNMP), command line interface (CLI), file transfer protocol (FTP), network time protocol (NTP), NetConf protocol entities and independent configuration files, so that different slices can also use the same IP address to manage devices.

– Control plane: Each slice has an independent internet group management protocol (IGMP), dynamic host configuration protocol (DHCP), dynamic routing protocols, etc.

– Forwarding plane: Each slice has an independent L2 and L3 forwarding engine and an independent virtual LAN (VLAN) and MAC table quota.

Compared to the traditional service isolation technology such as VLAN and virtual private network (VPN), slicing completely isolates the forwarding planes and control planes of each slice so that network security attacks and device faults in a slice do not affect other slices, such as MAC flooding attacks.

The OLT resources include physical and logical resources which can be allocated, reused and recycled on demand and exclusively occupied by slices to avoid resource collisions among slices. The physical resources include the uplink interface, PON board, PON interface, P2P interface and ONUs. The logical resources include the hardware forwarding table items such as VLAN, multicast, MAC, and ACL. The network identifier, such as VLAN ID or user IP address, can be reused in different slices with no conflict to facilitate the VNO's independent network planning.

ONU-level slicing scenario is shown in Figure 7-3, the bandwidth of PON interface is competed by all users under it which belongs to different VNOs. In the downstream direction, hierarchical QoS can be used for the bandwidth allocation, where the scheduler allocates the VNO's bandwidth at first according to the payments of different operators, and then further allocates bandwidth of each end-user within the VNO's capacity. In the upstream direction currently, the DBA scheme is based on the transmission container (T-CONT) that corresponds to one user or one service of a user. As there is no operator level bandwidth allocation, bandwidth of different VNOs cannot be well guaranteed and isolated. Upstream PON slicing is required in this case.

7.4 Use case 4: Slicing in PON systems for a multi-service shared network

Network slicing is used to divide a physical network into multiple logical devices which could reduce repeated network investments and implementations. It could also increase service convergence, security isolation and differentiated services. In PON systems, network slicing can be supported by slicing in the OLT, PON and ONU. By slicing a physical OLT CT into multiple PON slices, key resources such as access equipment rooms, OLTs, and ODNs can be shared. Differentiated services and independent operations and maintenance (O&M) requirements of different types of services and customers can also be realized with the PON slicing. PON slicing can be mainly used in two scenarios: slicing for multi-service shared networks and slicing for multi-operator shared networks. This clause provides an example of slicing for multi-service shared network in a PON system. Operators can use the same PON system to operate multiple types of services that provide differentiated service level agreements (SLAs). A typical multi-service scenario may include home broadband, and government and enterprise services. Even in home broadband service, users may have different applications and different applications may have different requirements.

This clause will introduce a use case of PON slicing for multi-services in detail. Figure 7-4 shows a typical architecture for PON slicing in this scenario. In order to meet different SLA requirements, the network resources can be divided into several different slices. Different slices may have different requirements and SLAs. Generally, the slices could be categorized into three different levels: golden slice with high priority, silver slice with middle priority and bronze slice with low priority. Different services are recognized by ONUs and are further assigned to specific slices based on their SLA requirements. The upstream PON slicing is realized by using QoS together with DBA, while the downstream PON slicing can be realized using QoS.

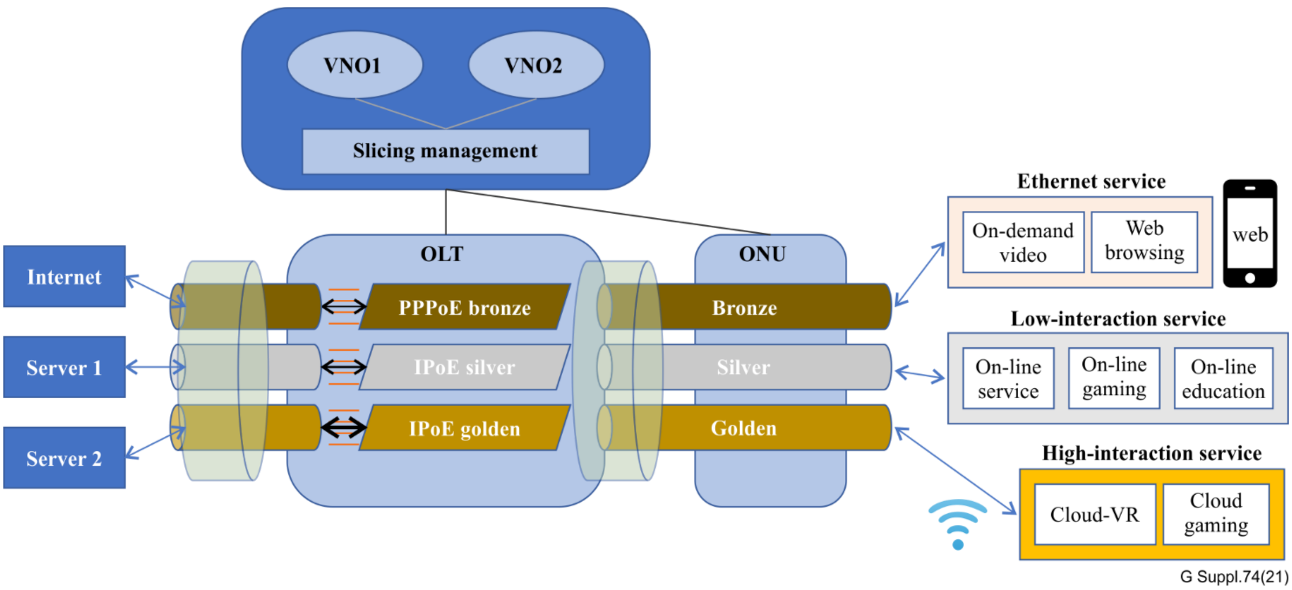

Figure 7-4 – Architecture of PON slicing for a multi-SLA

1) Golden slice (High priority)

The golden slice has the highest priority. It is always reserved for the operator's self-operating services with strong interaction which requires low latency and guaranteed bandwidth, such as cloud-VR and cloud gaming. In order to satisfy the bandwidth requirements, the bandwidth for golden slice should be assured and the bandwidth can be only shared within the golden slice. To satisfy the latency requirements, fixed bandwidth assignment shall be used and the QoS algorithm shall also be optimized.

In a golden slice, both bandwidth and latency should be guaranteed. One typical application of the golden slice is the cloud-VR. In general, the cloud VR requires bandwidth on the order of 100 MHz, the round-trip latency should be no larger than 20 ms. The time jitter tolerance is on the order 10 ms, and the packet loss rate should be less than 10^{-5}.

2) Silver slice (Middle priority)

The silver slice has a middle priority. It still has a high requirement for latency, while the bandwidth requirement is not quite high. In order to satisfy the requirement, the bandwidth for a silver slice is also guaranteed but can be shared with other slices. The silver slice can be used for some online services such as online gaming, online education and remote working.

For online gaming, the network mainly transmits the instruction content which requires low bandwidth, typically, 2 Mbit/s could be ideal. However, it has high requirements on latency and jitter. If the jitter exceeds 100 ms, frame freezing may occur, and if the jitter exceeds 250 ms, users' experience will be greatly affected, and the game cannot be played fairly. Table 7-1 shows the typical requirements for online gaming.

Table 7-1 – Network requirements of online gaming

Bandwidth	Average latency	Jitter	Package loss rate
2 Mbps	< 60 ms	< 100 ms	< 1.0E-1

Remote working services also require real-time performance because of nearly zero buffering. Therefore, there is a high requirement for time jittering. Jitter may cause frame freezing, frame skipping, and asynchronization of audio and video. For a typical case of 720p with a 15 fps frame rate, the frame interval is about 70 ms. If the delay jitter exceeds 200 ms, the experience will be affected. If 4K with 30 fps is used, the delay jitter tolerance shall be less than 100 ms. Table 7-2 shows the typical requirements for remote working services.

Table 7-2 – Network requirements of remote working services

	Bandwidth	Average latency	Jitter	Package loss rate
720p	5 Mbps	< 150 ms	< 200 ms	< 1.0E-2
4K	20 Mbps	< 60 ms	< 100 ms	< 1.0E-3

(3) Bronze slice (Low priority)

The bronze slice has an even lower priority. It may have a relatively high requirement for the bandwidth but a low requirement for the latency. The bandwidth type could be non-assured or best-effort. The bronze slice can be used for most of the typical Ethernet services such as on-demand video and web browsing. A large buffer could be used for such services, so that the experience is not sensitive to the latency and jitter. HDTV video could be a typical service with high bandwidth requirements. The bandwidth of a single 8K program after encoding ranges from 80 Mbit/s to 140 Mbit/s. Table 7-3 shows the typical requirements for 4K and 8K video services.

Table 7-3 – Network requirements of HDTV

Service type	Bandwidth	Average latency	Jitter	Package loss rate
4K video	~54 Mbps	/	/	< 1.0E-5
8K video	~150 Mbps	/	/	< 1.0E-6

Figure 7-4 shows that the slicing is based on the service flow. In the upstream, the service flow goes through the UNI of the ONU, ONU packet processing and forwarding, PON link transmission, OLT processing and forwarding, and the SNI interface. Achieving end-to-end slicing requires all the elements on the service flow path to support the slicing function, which implies the following supported functions in the ONU (1, 2) and OLT (3, 4, 5):

1) The ONU needs to identify application and service type and then distributes the packets to the corresponding slices.
2) The UNI at the ONU side shall support slice capability to ensure the SLA of different services.
3) The OLT needs to allocate the buffer and support forwarding and traffic management capabilities based on the PON slices, to guarantee different SLA levels.
4) The OLT needs to support packet forwarding for some slices and may also support TDM-like forwarding for some other slices to meet the requirements of deterministic low-latency services.
5) The SNI of the OLT needs to interconnect with the SNI of service nodes, refer to Table 7-4 below to achieve low latency and meet the SLA requirements.

Table 7-4 – Examples of SNI and services

SNI (Note 1)	**Physical interface (Note 2)**	**Service (Note 3)**
1GE/10GE/25GE/40GE/50GE/ 100GE/200GE/400GE – [b-IEEE 802.3]	1G/10G/25G/40G/50G/100G/ 200G/400G BASE	Ethernet, or Ethernet-based eCPRI (see [b-eCPRI])
[b-ITU-T G.703]	Plesiochronous digital hierarchy (PDH),	DS3, E1, E3, DS1, DS0
[b-ITU-T G.957]	STM-1, 4, 16, 64	E1, E3, DS1, DS3, GFP, E4, STM-n, DS0
[b-ATIS 0600107]	PDH	DS0, DS1, DS3
Synchronous digital hierarchy (SDH) / Synchronous optical networking (SONET)	SDH / SONET	OC3 – OC192
Optical transport networking (OTN) [b-ITU-T G.872 and b-ITU T G.709]	OTU1, OTU2, OTU3	OTN
CPRI / OBSAI (Open base station architecture initiative)	Option2, Option3 Option7, Option8	

Table 7-4 – Examples of SNI and services

NOTE 1 – There are many other services accommodated in high speed PON (HSP) systems, but these services do not have specified SNIs. NOTE 2 – Each item in the "Physical interface" column is illustrated by the corresponding entry in the "SNI" column. NOTE 3 – The column labelled "Service" shows which services can be supported by the physical interface.

7.5 Use case 5: PON slicing for industrial PON scenario

PON technology has become an advanced technology for factories, helping intra-plant network innovation. A typical implementation is to carry all the services from various sub-networks of the factory in a single PON system. Services of different sub-networks, such as the industrial process field data network, office network, surveillance network, are carried in the corresponding slices established in the OLT, as shown in Figure 7-5.

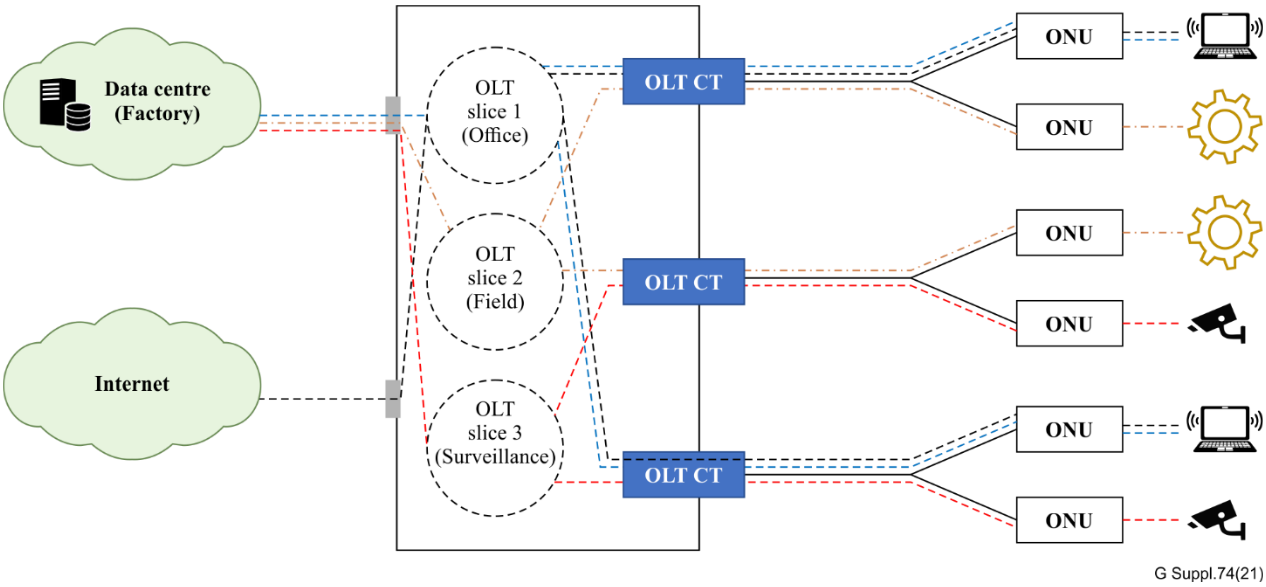

Figure 7-5 – Architecture of OLT slicing in an industrial PON scenario

Several slices (3, in the example in Figure 7-5) can be established in the PON system. This can be supported by OLT slicing, PON slicing and associated ONU slicing which are established for each OLT CT. Different slices may be mixed in the real deployment according to the factory demands. The ONUs that supports multiple slices may carry two or more services, each belonging to a corresponding slice with a specific SLA on parameters such as bandwidth, latency and jitter.

When employing slicing techniques in an industrial PON scenario, the services belonging to different slices need to be isolated with certain SLAs guaranteed in the forwarding plane. Independent management of different slices is better in order to minimize the interference effect. One important point is to make sure that the operation of one slice cannot impact other slices.

8 Functions needed for PON slicing in PON-based optical access networks

8.1 General

PON slicing involves the grouping of flows, for example, by grouping T-CONTs based on Alloc-ID. Slicing applies to both upstream and downstream and slices can be made identifiable in the access network (at the UNI/SNI).

Slicing should be work-conserving and, as far as possible, fully use the overall OLT CT bandwidth, although a part of the bandwidth might need to be reserved to make it assignable to different slices. Slices are typically assigned a minimum and a maximum bandwidth and unused bandwidth can be assigned to other slices if necessary. Slices can be assigned weights to apportion unused bandwidth.

PON already has mechanisms that can be reused for slicing such as T-CONTs and DBA in the upstream and traffic shaping mechanisms in the downstream.

The impact of slicing on ONU can vary. An ONU can be slicing-unaware if the applied slicing method is completely based on existing PON capabilities. In this case, an ONU might be dedicated to a known single slice type e.g., to optimize performance. A slicing-aware ONU might be designed to differentially treat multiple slices. ONU management might include identifying slices.

NOTE – OLT slicing and slicing-aware ONUs are for study in future versions of this Supplement.

8.1 Slice awareness in upstream

In general, bandwidth allocation to a slice or to a flow covers guaranteed and non-guaranteed bandwidth. The non-guaranteed bandwidth is allocated according to a fairness mechanism.

Slice awareness introduces a new level for the DBA, in addition to the existing flow level. The 2-level slice-aware DBA process can be described at a high level as follows:

1) Slice level

   1.1) Checking of slice saturation (done for all slices):

   - The demand of each flow inside the slice is estimated
   - The guaranteed bandwidth is assigned to every flow
   - Based on these assignments, each flow is determined to be either saturated or unsaturated
   - Based on these assignments, the slice is determined to be either saturated or unsaturated.

   1.2) Sharing of non-guaranteed BW between all unsaturated slices:

   - The surplus bandwidth on the channel termination is calculated based on the total capacity and the allocated guaranteed bandwidths in step 1.1.
   - The surplus bandwidth on the channel termination is then shared between the unsaturated slices according to a fairness criterion.

2) Flow level – sharing of non-guaranteed BW between all unsaturated flows within a slice (done per slice, for all slices):

   - The process is identical to the PON TC Recommendations, but within the constraints of a given slice: the surplus bandwidth allocated to the slice in step 1.2 is shared between the unsaturated flows of the slice according to a fairness criterion.

The fairness criterion at flow level can be based on flow-level rate proportionality or on flow-level priority and weight (as defined in the PON TC Recommendations, e.g., clauses C.7.3.5 and C.7.3.6 in [ITU-T G.9807.1]).

Similarly, the fairness criterion at the slice level can be based on the slice-level rate proportionality or on the slice-level priority and weight.

A slice-aware DBA provides the following:

- the possibility to configure upstream service parameters at slice level.
- the possibility to assign guaranteed and non-guaranteed bandwidth at the slice level, in an efficient way (avoiding wasted upstream capacity). The non-guaranteed bandwidth is allocated to the slices according to a fairness criterion.
- the option to use bandwidth rates configured at slice level for the slice-level rate-based fairness criterion between unsaturated slices is:
  - The surplus bandwidth is exhausted and at most one slice remains unsaturated, or
  - The surplus bandwidth is exhausted and at least two slices remain unsaturated, where for any two unsaturated slices their assigned non-guaranteed bandwidths satisfy a rate-based fairness condition.
- the option to use priority level and a weight configured at slice level for the fairness criterion between unsaturated slices is:
  - As long as at least one slice with a provisioned priority level remains unsaturated, the assigned non-guaranteed bandwidth share of any slice with a logically lower provisioned priority level is zero.
  - As long as at least two slices with identical provisioned priority levels remain unsaturated, their assigned non-guaranteed bandwidths satisfy a weight-based fairness condition.
- the possibility to provide isolation between slices on the CT for their guaranteed characteristics. Isolation can refer to one or both of the following characteristics:
  - Capability to allocate a guaranteed upstream bandwidth to a slice, that is always obtainable for the slice, independent of the bandwidth allocated to and consumed by the other slices, within the limit of the total capacity of the considered channel termination.
  - Capability to control upstream timing properties such as time jitter tolerance (T_{JT}), bandwidth assignment delay tolerance (T_{BDT}), protection switching delay tolerance (T_{PST}) of flows inside a slice, independent of bandwidth allocations to other slices, within the limit of the bandwidth map (BWmap) generation capabilities per PON frame.

    NOTE – These timing properties are defined in the existing PON TC Recommendations and are the properties of a flow, not of a slice. For example, a slice does not have a T_{JT} on itself, but every flow in the slice can be configured with a T_{JT}, possibly with slice-specific constraints on its value (min, max).

8.2 Slice awareness in downstream

In general, bandwidth allocation to a slice or a flow covers guaranteed and non-guaranteed bandwidth. The non-guaranteed bandwidth is allocated according to a fairness mechanism.

Note that the QoS-aware traffic management of a downstream flow to the PON channel termination is not described in the PON TC Recommendations as it can be done locally at the channel termination in the OLT, independently of the ONUs and is implementation-specific. Similarly, the traffic management of downstream slices is not technically described in this Supplement.

This clause identifies the set of functional expectations on the downstream traffic management at the channel termination for the handling of downstream slices:

- the possibility to configure downstream service parameters at slice level.
- the presence of a slice level in the downstream (hierarchical) scheduling.
- the possibility to assign guaranteed and non-guaranteed bandwidth at the slice level, in an efficient way (avoiding wasted downstream capacity). The non-guaranteed bandwidth is allocated to slices according to a fairness criterion.

- the possibility to provide isolation between slices on the CT for their guaranteed characteristics.

8.3 Control of the common per-PON CT resources

The PON system needs to preserve the guaranteed performances of existing slices and existing flows in a live network from changes in the number of slices, or from conflicts with additional per-slice or per-flow configurations.

Therefore, a control mechanism is needed during configuration to ensure that any changes in the configuration will not have negative impacts on the performances that are guaranteed, in particular at the level of an OLT CT.

The check to be performed at each new configuration would determine whether the associated guaranteed parameters of the requested configuration are not in conflict with the remaining capacity on the OLT CT, in terms of bandwidth and timing parameters.

Such a check should be done at the corresponding level:

– at OLT CT-level when adding a new slice (check on guaranteed bandwidth).

– at OLT CT-level when adding a new flow in a slice (timing property).

– at slice level when adding a new flow in the slice (check on guaranteed bandwidth).

When the new configuration passes the checks, it can be safely activated on the PON system without affecting the guaranteed parameters of existing and active slices and flows. The allocation of non-guaranteed parameters will definitely be impacted by new configurations as the common capacity (e.g., surplus bandwidth) capacity will be shared by more entities.

To start using an OLT CT without slicing and activate slicing at a later stage without impacting the guaranteed properties of the existing flows, it is recommended that the operator first plans and reserves some capacity on the OLT CT for future slice guarantees. The OLT CT would start with all flows being grouped in a single default (or "null") slice with the following characteristics:

– the guaranteed bandwidth of the default slice can be as high as the left-over OLT CT capacity (i.e., total CT capacity – reserved guarantees for future slices), or lower.

– the maximum of the default slice can be as high as the full OLT CT capacity or lower.

When slicing is then introduced on the OLT CT at a later stage, new slices are created. The same checks as mentioned above will allow preserving the guarantees of the default slice and its flows. Each newly added flow is then associated with the default slice or with a newly created slice.

NOTE – If an ONU is slice-aware, its re-configuration may be needed when changing the slice-level characteristics of (one of) its slice(s), or when changing the association(s) of such an ONU to given slice(s) (either by deleting or adding slice(s)).

8.4 Functions related to the scope of a PON slice control

Slicing can be used by a network operator in two general styles, either to partition its own services or to share its network with other operators. A network operator that partitions its own services across the PON has visibility and control over all slices and will configure slice bandwidth and other parameters to match the needs of each slice. A network operator using a shared PON is likely to have visibility and control only of the slices assigned to them, with overall PON control handled at a higher level. These different styles imply two different sets of behaviour for assigning PON CT bandwidth to slices, as follows:

For a slice user with control only over the slices assigned to them, a slice needs to behave as though it is dedicated to that user, including:

- bandwidth is always available

- the slice user may need bandwidth to be available without re-assignment delays, although a delay may be acceptable depending on a service level agreement
- a slice that is not carrying any traffic can safely be considered idle, for example, ONUs that are on that slice only can be put into a low power state
- bandwidth is capped with an upper limit, even if spare bandwidth exists on the PON CT.

For a slice user with control over all slices:

- guaranteed bandwidth is always available, potentially with re-assignment delays
- slice bandwidth that is not used is considered available for re-assignment to other slices
- bandwidth has no fixed upper limit other than the total bandwidth of the PON CT.

Note that this behaviour is as seen by the slice user and need not constrain how they are met by the system. For example, if re-assignment of in-use bandwidth takes the same time as an assignment of unused bandwidth, the same procedure could be used for both types.

9 Using slice-aware DBA

9.1 Dynamic bandwidth assignment overview

As defined in the PON transmission convergence specifications like [ITU-T G.9804.2], dynamic bandwidth assignment (DBA) is the process by which the OLT allocates upstream transmission opportunities to the traffic-bearing entities within the ONUs of a given PON channel, based on the dynamic indication of their activity and their configured traffic contracts. The activity status indication can either be explicit through buffer status reporting, implicit through the transmission of idle XGEM frames during the upstream transmission opportunities, or both.

In comparison with static bandwidth assignment, the DBA mechanism improves upstream bandwidth utilization by reacting adaptively to the ONUs' burst traffic patterns. The practical benefits of DBA are twofold. First, the network operator can add more subscribers to the access network due to more efficient bandwidth use. Second, subscribers can enjoy enhanced services, such as those requiring variable rates with peaks extending beyond the levels that can reasonably be allocated statically.

DBA can either be performed solely at the level of the individual traffic-bearing entities (flows), or it can optionally be performed at two levels: namely at flow level and additionally at slice level by taking groups of individual traffic-bearing entities into account. Each group represents a slice, and, in that case, DBA is applied with awareness of the multiple slices. In the context of PON virtualization, when providing a logical network using resources (a virtual PON) to support various services, slice-aware DBA can provide an optimal use of upstream PON resources while maintaining service quality.

9.2 Slice-aware DBA

The slice-aware DBA can secure isolation between logical networks and can operate stable services with optimal resources within each PON slice.

Slice-awareness is provided by grouping all T-CONTs based on the Alloc-ID into slices with an associated slice-level traffic descriptor.

The DBA then assigns a slice-level bandwidth for each slice. There are two ways to set such bandwidth boundary of a PON slice, namely statically or dynamically. In the static approach, the assigned bandwidth to a PON slice is a fixed value isolating it from the other slices. With a dynamic approach, bandwidth-isolation between PON slices is performed dynamically based on a multi-slice DBA, assigning per slice bandwidths (whereby each slice is constrained by a maximum bandwidth).

The DBA then assigns flow-level bandwidth for each T-CONT inside a slice and does this for each slice. An ONU can belong to multiple PON slices at the same time by using multiple T-CONTs pertaining to different slices. Note that when slicing is used, every T-CONT is part of exactly one slice. This is illustrated in Figure 9-1.

In both cases, the PON slice is configured and monitored by a control architecture as shown in clause 10.

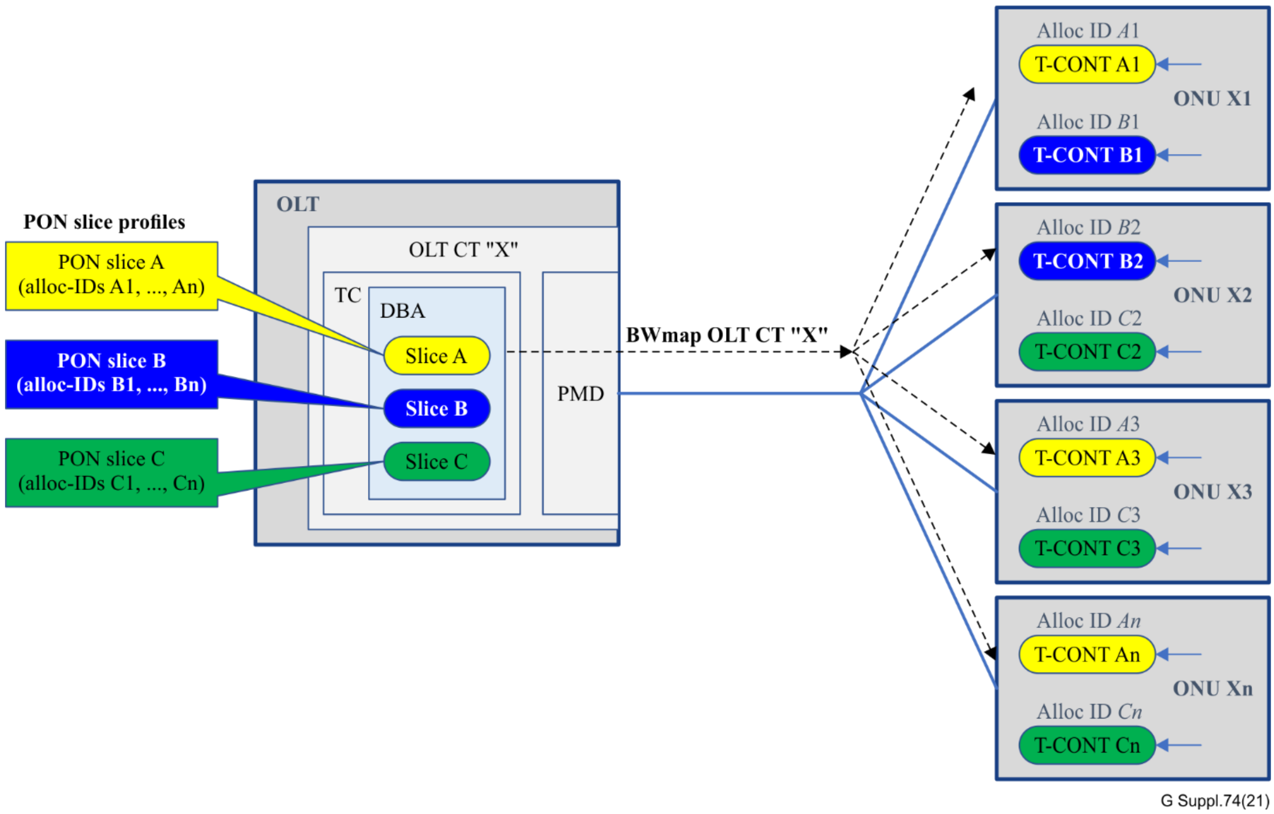

Figure 9-1 – Slice-aware DBA per PON CT (Abstract view on PON CT "X" with slices A, B, C)

Creating a BWmap can be similar to the "HSP channel DBA abstraction" [ITU-T G.9804.2], but the part to be calculated is different for each slice. The slice-aware DBA function calculates on each slice with different properties, combines the results, and creates a BWmap.

The BWmap compositing method involves allocating occupied space (cycle, bandwidth) to meet the properties of each slice.

Figure 9-2 shows an example of a BWmap combining bandwidth allocations for each slice. The BWmap is separated into regularly occurring blocks of contiguous allocations for each slice. Blocks for low latency services are allocated more frequently to ensure latency properties.

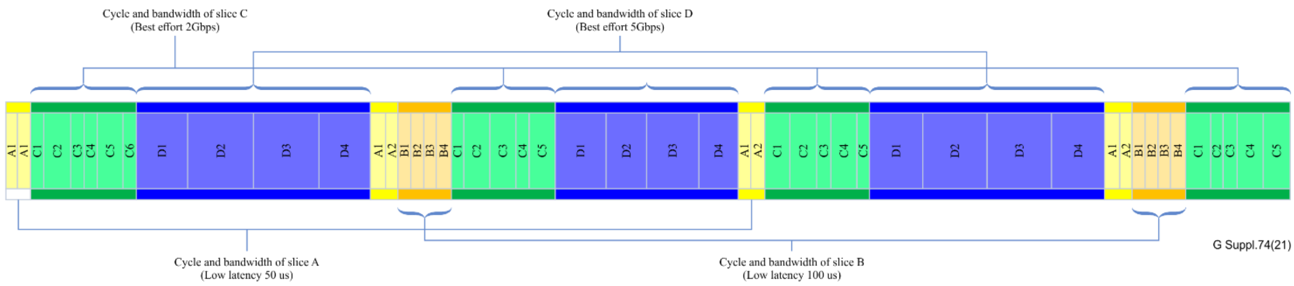

Figure 9-2 – BWmap example with combining bandwidth allocation for each slice

More sophisticated approaches can be considered with the interleaving of several sets of allocations of different slices to provide lower latency for instance.

BWmap approaches can also optimize the PON resource while guaranteeing the desired level of isolation between slices.

10 Control of PON slices

10.1 Introduction

Network slicing contributes to and integrates into global telecommunications systems. To realise networks in software, control of network slicing is expected to use a software defined network (SDN) as described in several ITU Recommendations ([ITU-T Y.3150], [ITU-T Y.3151]) and Broadband Forum technical report [BBF TR-370] on fixed access network sharing. The PON management system needs to support the necessary slice-level configurations to allow the behaviour of slicing at the CT level described in clause 8.

10.2 Architecture description

This clause shows generically how the PON (passive optical network) equipment relates to the architecture supporting network slicing. Figure 10-1 can be specifically related to [ITU-T Y.3151] Figure 6-1 in the SDN context for an example.

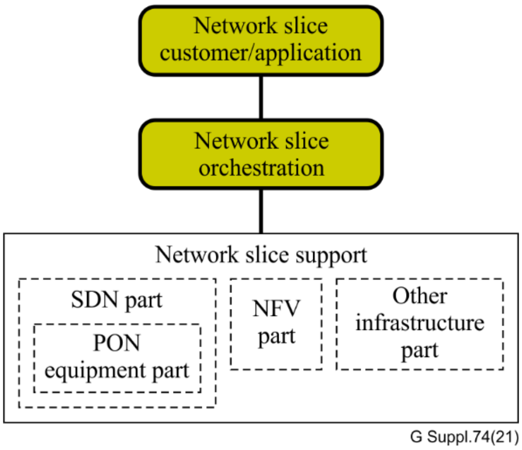

Figure 10-1 – High-level architectural model for network slice support in PON

- Architecture model for PON slice

The SDN part of the network slice support consists of an SDN controller and an SDN infrastructure. Figure 10-2 illustrates the SDN relevant functional blocks within the SDN part in Figure 10-2. The basic functional roles of the SDN controller and the SDN infrastructure are as follows:

– SDN controller provides virtual networks (VNs) composed of various PON slices according to the requests of the network slice orchestration. A PON slice is part of the virtualized resources (VRs) in [ITU-T Y.3151].

– SDN infrastructure (PON systems): provides PON slices as virtualized resources (VRs) according to the requests of the SDN controller.

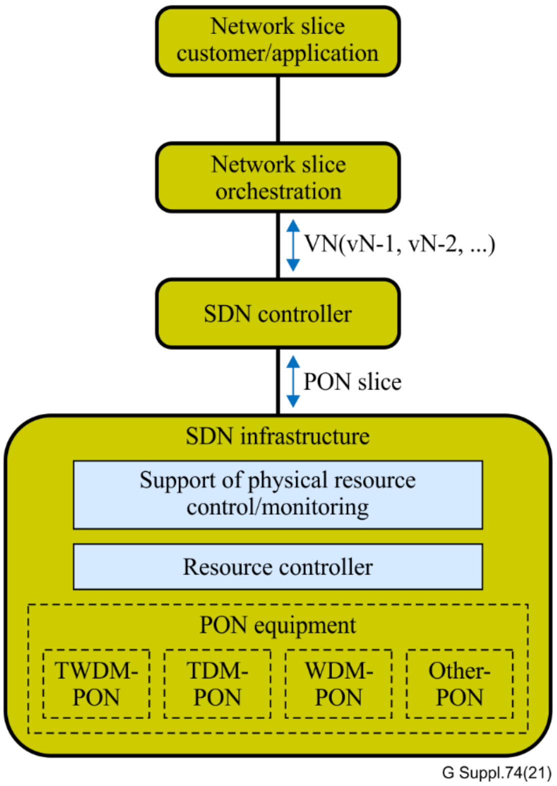

Figure 10-2 – PON system architectural model of SDN part

The SDN infrastructure consists of a resource controller and various PON equipment. The role of the resource controller is to control the OLTs. There are varieties of OLT hardware and various possible methods to implement this control in practice. Examples of PON equipment are time division multiplexing (TDM-PON), wavelength division multiplexing (WDM-PON), time and wavelength division multiplexed (TWDM-PON), etc. The SDN infrastructure has functions supporting the monitoring and controlling of the PON equipment.

- Management function

Management aspects of the SDN controller are as follows:

– Management aspects for the SDN controller are described in [ITU-T G.7702].

– An element management system (EMS) needs to send and receive management signals for interactions with the operating support system (OSS).

10.3 Control interface

This clause describes the generic SDN control interface to handle network slicing and its application to PON.

The control interface of network slice is defined using reference points and related parameters.

- SDN reference point

The reference point relevant to SDN is R2S is shown in Figure 10-3:

– R2S: a control interface between the SDN controller and an SDN infrastructure (PON systems). R2S refers to and expands the standardized specifications of [ITU-T G.7702].

– The abstraction of the SDN controller converts between virtual networks (VNs) and various PON slices (with different properties, bandwidth, latency, etc.) through R2S. When configuring a virtual network (VN) directly with PON slices, SDN abstraction is not needed.

– The PON abstraction of the resource controller in an SDN infrastructure converts between PON slices and various PON equipment (e.g., G-PON, XGS-PON, TWDM-PON) through PON-specific interfaces.

– The resource controller accounts for functions that apply on PON channel termination ports and can be performed for each PON slice with separate parameters. The slice-aware DBA is an example of these functions.

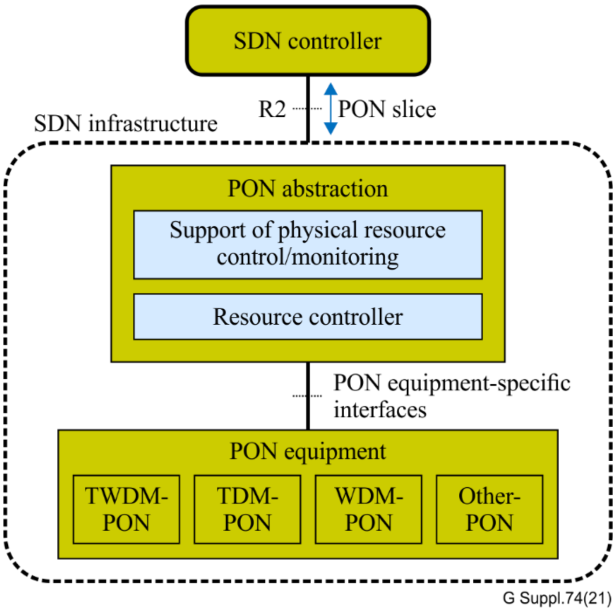

Figure 10-3 – Reference points of SDN for network slicing in PON

• Types of parameters

Parameters transmitted via R2S consist of general parameters based on the specifications of [ITU-T G.7702] and expanded parameters for create, modify, update and delete (CMUD) of PON slice.

The general parameters are defined in [ITU-T G.7702] and [ITU-T Y.3151].

Expanded request parameters which are sent from an SDN controller to an SDN infrastructure (PON systems), can be used for getting information of the SDN infrastructure (PON systems) and for allocation of resources for PON slices in the SDN infrastructure (PON systems).

Expanded result parameters which returns from the SDN infrastructure (PON systems) to the SDN controller can be used for replying to information of the SDN infrastructure and for monitoring the status of the PON slices in the SDN infrastructure.

• Control interface between SDN controller and PON systems

1) Information for PON slice creation, monitor, update and delete (CMUD)

Information between the SDN controller and the PON systems using the expanded request parameters and expanded result parameters are used for configuration, resource allocation and resource monitoring of PON slices. Note that the specific parameters corresponding to these information elements depend on the degree of abstraction in the PON abstraction function.

– Configuration: number of ONUs, number of logical links, data flow identification, logical port (virtualized UNI/SNI) bandwidth, latency (class), others are for further study.

– Resource allocation: QoS parameters, per slice traffic descriptor, others are for further study.

– Resource monitoring: Physical resource status (usage, performance, etc.).

2) Details of the interface (Message format, protocols, sequence of process, etc.).

For further study.

Bibliography

[b-ITU-T G.703] Recommendation ITU-T G.703 (2016), *Physical/electrical characteristics of hierarchical digital interfaces*.

[b-ITU-T G.709] Recommendation ITU-T G.709/Y.1331 (2020), *Interfaces for the optical transport network.*

[b-ITU-T G.872] Recommendation ITU-T G.872 (2019), *Architecture or the optical transport network.*

[b-ITU-T G.957] Recommendation ITU-T G.957 (2006), *Optical interfaces for equipments and systems relating to the synchronous digital hierarchy*.

[b-ATIS 0600107] ATIS 0600107 (2005), *Digital Hierarchy – Formats and Specifications.* <https://standards.globalspec.com/std/9917437/atis-0600107>

[b-eCPRI] eCPRI Specification V2.0 (2019), *Common Public Radio Interface: eCPRI Interface Specification.* <http://www.cpri.info/downloads/eCPRI_v_2.0_2019_05_10c.pdf>

[b-IEEE 802.3] IEEE 802.3 (2021), *Ethernet working group*. <https://www.ieee802.org/3/>

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	Tariff and accounting principles and international telecommunication/ICT economic and policy issues
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	**Transmission systems and media, digital systems and networks**
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Environment and ICTs, climate change, e-waste, energy efficiency; construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling, and associated measurements and tests
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	Global information infrastructure, Internet protocol aspects, next-generation networks, Internet of Things and smart cities
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2022